arXiv:astro-ph/0511414v2 17 Jan 2006

CLUES TO AGN GROWTH FROM OPTICALLY VARIABLE OBJECTS
IN THE HUBBLE ULTRA DEEP FIELD

S. H. Cohen, R. E. Ryan Jr., A. N. Straughn, N. P. Hathi, R. A. Windhorst
Department of Physics and Astronomy, Arizona State University, Tempe, AZ 85287-1504

AND

A. M. Koekemoer, N. Pirzkal, C. Xu, B. Mobasher, S. Malhotra, L.-G. Strolger, J. E. Rhoads
Space Telescope Science Institute, Baltimore, MD 21218
Draft version February 5, 2008

ABSTRACT

We present a photometric search for objects with point-source components that are optically variable on timescales of weeks–months in the Hubble Ultra Deep Field (HUDF) to $i'_{AB} = 28.0$ mag. The data are split into four sub-stacks of approximately equal exposure times. Objects exhibiting the signature of optical variability are selected by studying the photometric error distribution between the four different epochs, and selecting 622 candidates as 3.0σ outliers from the original catalog of 4644 objects. Of these, 45 are visually confirmed as free of contamination from close neighbors or various types of image defects. Four lie within the positional error boxes of Chandra X-ray sources, and two of these are spectroscopically confirmed AGN. The photometric redshift distribution of the selected variable sample is compared to that of field galaxies, and we find that a constant fraction of $\sim 1\%$ of all field objects show variability over the range of $0.1 \lesssim z \lesssim 4.5$. Combined with other recent HUDF results, as well as those of recent state-of-the-art numerical simulations, we discuss a potential link between the hierarchical merging of galaxies and the growth of AGN.

Subject headings: Galaxies: active — galaxies: formation — Active Galactic Nuclei — Supermassive Black Holes

1. INTRODUCTION

The Hubble Ultra Deep Field (HUDF; Beckwith et al. 2005) is the deepest optical image of a slice of the Universe ever observed. As such, it allows for a variety of different investigations into astrophysical objects and processes. The HUDF observations consist of 400 orbits with the Hubble Space Telescope (HST) over a period of four months in four optical bands ($BVi'z'$) with the Advanced Camera for Surveys (ACS), and are supplemented in the JH bands with an HST Legacy Program using the Near-Infrared Camera and Multi-Object Spectrograph (NICMOS; Bouwens et al. 2004; Thompson et al. 2005). Given the unprecedented quality of these data, the list of supporting data from both space and the ground is constantly growing.

Since the ACS data was observed over a period of four months, it provides a unique opportunity to search for variability in all types of objects to very faint flux levels (perhaps even for $AB \gtrsim 30$ mag), such as faint stars, distant supernovae (SNe), and weak active galactic nuclei (AGN). In this paper, we search for weak AGN variability in the i'-band (F775W), because, in this filter, the HUDF images are deepest, and have the best temporal spacing over the four months for the variability search. It should be noted that at higher redshifts, the ACS filters sample further into the rest-frame ultraviolet. This is advantageous, because AGNs are known to show more variability in the UV (e.g., Paltani & Courvoisier 1994). In the original Hubble Deep Field North (HDF–N) and the Groth Strip Survey, Sarajedini et al. (2003) and Sarajedini (2003) performed a similar search to AB=27.0 mag over longer time baselines (5–7 years).

They searched for variability in the nuclei of the galaxies using small apertures, which is a different approach than in the present search.

From the WMAP polarization results (Kogut et al. 2003), population III stars likely existed at $z \simeq 20$. These massive stars ($\gtrsim 250\ M_\odot$) are expected to produce a large population of massive black holes ($M_{bh} \gtrsim 150\ M_\odot$, Madau & Rees (2001)). Since there is now good dynamical evidence for the existence of supermassive ($M_{bh} \simeq 10^6 – 10^9\ M_\odot$) black holes (SMBH) in the centers of galaxies at $z \simeq 0$ (Kormendy & Richstone 1995; Magorrian, Tremaine, & Richstone 1998; Kormendy & Gebhardt 2001), it is important to understand if there is any relationship between the formation of the SMBHs observed at $z \simeq 0$ and the lower mass BHs at z$\simeq$20. A comprehensive review of SMBHs is given by Ferrarese & Ford (2004). An important question to address is how these SMBHs, as seen nearby, have grown over the course of cosmic time. One suggestion is that they "grow" through the mergers of galaxies that themselves contain less massive BHs, so the byproduct is a larger single galaxy with–eventually–a more massive BH in its center. The growth of the BH may then be observed via its AGN activity. If this scenario is true, then perhaps there exists an observable link between galaxy mergers and increased AGN activity (Silk & Rees 1998). Therefore, studying this possible link as a function of redshift could give insight into the growth of SMBHs and its potential relation to the process of galaxy assembly.

In §2, we present the HUDF observations and summarize the essential elements of its data reduction, and in §3 we present the variable candidate selection. In §4 we present the photometric redshift distribution of the

Electronic address: seth.cohen@asu.edu

Epoch No.	Start Date	End Date	Exp. Time[a] (Seconds)	# of Exp.	RJD[b] (Days)	Days Since Epoch 1
1	2003-09-24	2003-10-10	92340	76	52914.2	⋯
2	2003-10-10	2003-10-29	92340	76	52926.7	12.5
3	2003-12-04	2003-12-18	89940	76	52985.9	71.7
4	2003-12-22	2004-01-14	72490	60	53005.7	91.5

[a]There are two exposures per HST orbit

[b]RJD=median Revised Julian Date-2.4×10^6 days

variable objects together with that of the HUDF field galaxies, in §5 we present the results, and in §6 we discuss our results in terms of galaxy assembly and AGN growth.

2. OBSERVATIONS

All data used here are from the Hubble Ultra Deep Field (HUDF; Beckwith et al. 2005). The individual cosmic-ray (CR) clipped images and weight maps were used with *multidrizzle* (Koekemoer et al. 2002) to create eight sub-stacks of approximately equal exposure times. These used the same cosmic-ray maps and weight maps employed to create the full-depth HUDF mosaics. All HUDF images were *drizzled* onto the *same* output pixel scale (0″030 per pixel) and WCS frame as the original HUDF. Given the time-spacing of the exposures and the desire to extend the study to the faintest possible flux-levels, the images and weight maps were combined in groups of two to create four epochs of observation for the variability study on 0.4–3.5 months timescales. Exposure–time weighted averages were created for all images, and simple addition was used to combine the weight maps (Beckwith et al. 2005). The exposure times and median observation dates are listed in Table 1. The four epochs chosen here have exposure times as close to each other as possible, so that the flux error distributions will be as much as possible symmetric, and therefore more easily modeled. As seen in Table 1, this is done at the expense of not having the endpoints of the four epochs well-spaced in time. In order to be optimized for variability studies, future observations of this kind should take into account the need for both equal depth exposures and well separated observation dates, although this would further complicate the already difficult task of scheduling observations such as the HUDF. All magnitudes are on the *AB*-system using the zero-points given in the HUDF public data release (Beckwith et al. 2005).

2.1. *Catalog Generation and Photometry*

Variability was searched for by comparing the photometric catalogs from the various epochs to each other. Catalogs were generated using *SExtractor Version 2.2.2* (Bertin & Arnouts 1996) with a 1.0σ detection threshold, and requiring a minimum of 15 connected pixels (i.e. approximately the PSF area) at this limit above sky. Since we are searching for any signs of variability, we chose to use a liberal amount of de-blending ($DEBLEND\_MINCONT = 5\times10^{-6}$). This allows for pieces of merging galaxies to be measured separately to enhance the chances of finding variable events in point-source components. *SExtractor* was run in dual image mode using the full-depth HUDF as the detection image,

and utilizing the corresponding weight-maps to minimize the number of false detections due to edges and other image defects that are reflected in these weight-maps. This results in catalogs with 27819 measured objects, which still contains many over-deblended objects or edge-effects. Since the HUDF was observed at four different position angles (Beckwith et al. 2005), only the 15205 objects observed in all four epochs were considered. The result is a catalog of 12514 objects, which is 90% complete to $i' \lesssim 30.5$ mag. Since we can only measure variability from the individual epoch images that are one-fourth the full HUDF in length, the variable candidate sample is restricted to the 4644 objects with $i' < 28.0$ mag.

The HUDF is the deepest optical image ever observed, and will possibly remain the deepest until the James Webb Space Telescope (JWST) is launched in 2011. To explore the limits of the HUDF depth, a few words about the point spread function (PSF) of the ACS images are needed. A comprehensive study of HST/ACS PSF-issues can be found in Heymans et al. (2004), so here we only highlight the aspects relevant to our study. First, the ACS camera is not located on the optical axis of HST, and therefore the HUDF field is rectified by applying geometric distortion correction polynomials. Secondly, the ACS/WFC PSF is known to vary with location on the CCD detectors, and with the time of observation due to "breathing" of the HST Optical Telescope Assembly (OTA). Since the HUDF was observed at four different roll angles over four months, these PSF effects can easily be seen by inspecting the locations of bright objects in an image created by dividing two images taken at different roll angles. Owing to these PSF issues and the significantly complex ACS image registration, this "ratio-image" will easily show the cores of all bright objects with significant positive or negative flux excursions, regardless of whether or not they are truly variable. For this reason, we *cannot* use small PSF-sized apertures to search for nuclear variability, as was done by Sarajedini et al. (2003). Sarajedini et al. (2003) *could* use the small–aperture method, because for the much larger WFPC2 pixel-size and the *on-axis* location of the WFPC2 camera, the geometrical distortion correction and registration effects are much smaller. Instead, we chose to use total magnitudes of the highly deblended objects. Even though our total flux apertures may encompass the whole galaxy, the variability necessarily must come from a small region (less than the 0″084 PSF), due to the finite light-travel time across the region of variability.

3. CANDIDATE SELECTION

The catalogs of the four epochs were all compared to each other resulting in six sets of diagrams similar to the one shown in Fig. 1. These figures show the change in measured magnitudes in the SExtractor matched apertures as a function of full-depth HUDF flux. In order to determine which objects were true outliers (i.e., variable candidates), the expected error distribution for each set was computed as follows. For each measurement in a given epoch, we compute the total flux error for that i^{th} flux measurement:

$$\sigma_i^{tot} = \sqrt{\frac{\sigma_i^2 A(F_i) + F_i/G}{F_i^2}} \qquad (1)$$
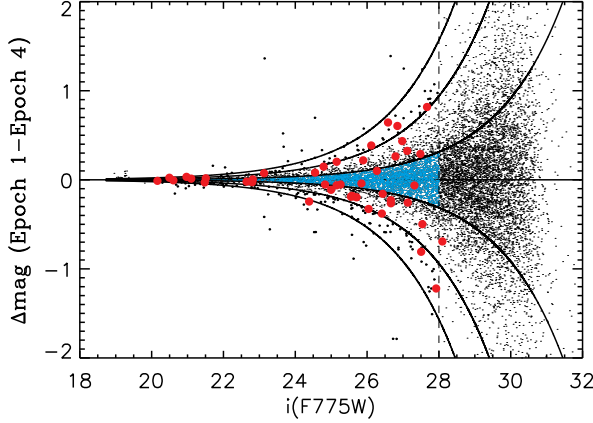
FIG. 1.— Magnitude difference between two HUDF epochs of all objects versus their total i'-band flux in AB-mag from matched total apertures. The $\pm 1\sigma, \pm 3\sigma, \pm 5\sigma$ lines are shown. The black circles show the 222 variable candidates chosen to be $|\Delta mag| \gtrsim 3\sigma$ outliers between any one of these these two epochs and $20 \lesssim i' \lesssim 28.0$ mag. Blue points show the $|\Delta mag| \leq \pm 1\sigma$ points used to normalize the error distribution, so that 1.0σ reflects as much as possible the true Gaussian 1-σ line. Large red points show the "best" 45 candidates that were chosen from all six possible epoch combinations, many of which were seen at $\gtrsim 3.0\sigma$ in 2 or more epoch combinations.

where σ_i is the RMS per pixel in the sky background, $A(F_i)$ is the number of pixels belonging to object i of a given flux (described in detail below), F_i is its measured flux in e^- per second and G is the gain in units of seconds. This quantity is computed for each epoch, and for each of the epoch-pairs they are combined in quadrature:

$$\Delta mag = \pm \sqrt{(1.0857N)^2 \times \left((\sigma_i^{tot})^2 + (\sigma_j^{tot})^2\right)} \quad (2)$$

where N is the number of σ for which the error-model is computed, and i, j denote the measurements at a given flux-level in each of the two epochs under consideration. Since each object at a given flux level can have a different area A (i.e., number of pixels), we need to assume a general relation between flux and area in order to optimally model the true error distribution using the above equations. This relation was determined iteratively for each pair of observations, such that 68.3% of the points lie within the boundaries of the upper and lower 1.0σ lines. We started this process by fitting the relation between the *SExtractor* magnitude and *ISOAREA_IMAGE* parameters as a first guess. It is then assumed that flux is proportional to area, and we solve iteratively for the proportionality constant to get $\pm 1.0\sigma$ lines that maximally represent a Gaussian error distribution (Fig. 1). In order to demostrate the Gaussian nature of this error distribution at all flux levels, the Δmag data are divided by the $1.0-\sigma$ model line, and histograms were computed for the resulting Δmag data at various flux-levels (Fig. 2). These histograms are remarkably well fit by Gaussians of $\sigma \simeq 1$. The HUDF noise distribution is not perfectly Gaussian, but with 288 independent exposures in the i'-band, the error distribution is as close to Gaussian as seen in any astronomical CCD application.

Once this 1σ line is determined, we set $N = 3.0$ in Eqn. 2, and find all the objects that are at least 3.0σ outliers. In Fig. 1, we show the $\pm 1\sigma$, $\pm 3\sigma$, and $\pm 5\sigma$
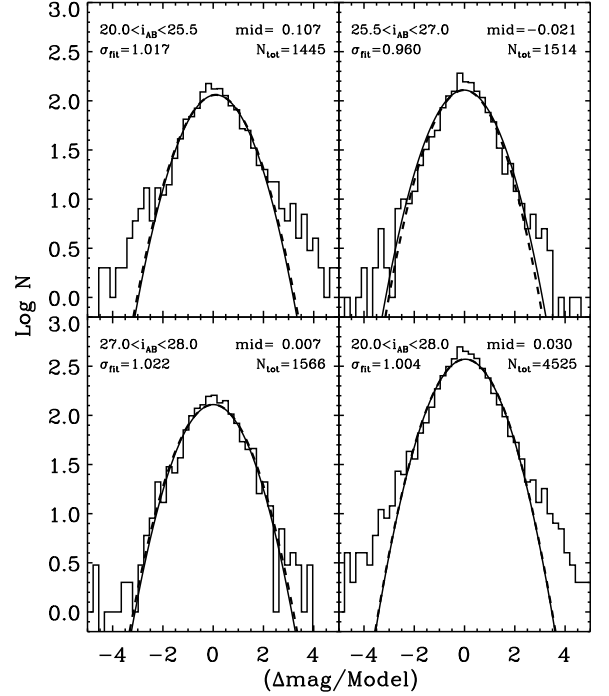

FIG. 2.— Gaussian nature of the flux error distribution at all flux levels. The Δmag data from Fig. 1 are divided by the best–fit model 1.0-σ lines. Histograms are computed from the resulting data for the indicated magnitude ranges. All distributions are well fit by Gaussians (parabolas in log space) with $\sigma \simeq 1.0$ as indicated in the individual panels. The almost indistinguishable dashed and solid lines are for the best–fit and $\sigma \equiv 1$ Gaussians, respectively.

lines, along with the actual data. For this particular pair of catalogs, there were initially 222 objects which varied in flux by more than 3.0σ. The choice of 3.0σ implies that we can expect 0.27% random contaminants. Among 4644 objects to $i' = 28.0$ mag, this is about 13 interloping objects that are potentially classified as variable, simply because of the chosen 3.0σ cut-off.

4. PHOTOMETRIC REDSHIFTS

Photometric redshifts were computed in two ways. In the first method, magnitudes are computed in a fixed aperture with a one arcsecond radius, and objects are selected from the ACS z'-band. In addition to the HUDF ACS $BVi'z'$ images, the available NICMOS JH (Thompson et al. 2005) and VLT ISAAC K band images are used where available. The z'-band selection allows for $z \gtrsim 5.5$ galaxies to be included in the list, but the z'-band is not as deep as the i'-band or V-band images in the HUDF. Therefore, the primary object definition catalog was made in the i'-band, which introduces a bias against $z \gtrsim 5.5$ objects (see Yan & Windhorst 2004). The use of the large apertures allows for the ground-based seeing K-band fluxes to be included for more precise photo-z estimates. One possible problem is that most faint galaxies are significantly smaller than these apertures, such that problems may arise for objects in crowded regions.

In an attempt to address these issues, we also tried using magnitudes measured within apertures defined in the i'-band, using the same apertures in which the variability was searched for (*SExtractor* parameter *MAG_AUTO*). The i'-band selection limits us to objects with $z \lesssim 5.5$.

This photometry is only applied to the ACS $BVi'z'$ and NICMOS JH data, which have the necessary resolution to accurately measure fluxes on sub-arcsecond scales. The VLT K-band data are not used here, because they are limited by ground-based seeing ($\lesssim 1$" FWHM). The disadvantage of not using the K-band is the lower redshift accuracy, but the advantage is that the flux is measured from the *same* object component in all filters, so that crowding is less of an issue for this method. The fluxes and errors measured in this way are then input into the photometric redshift code *hyper-z* (Bolzonella et al. 2000), using a suite of both empirical (Coleman, Wu, & Weedman 1980) and evolutionary spectral synthesis (GISSEL98 update to Bruzual & Charlot 1993) templates.

While these methods each have their own advantages and disadvantages, they produced the same important results discussed below. The difference between the methods are minor in the photometric redshift distribution produced, and since we discuss ratios of photometric redshift distributions in what follows, these differences are not relevant for the main argument below. The second photometric redshift determination method was adopted for all figures shown here.

5. RESULTS

5.1. *Number of Variable Objects*

We find unique 622 out of 4644 variable candidates with $i' \leq 28.0$ mag from the six possible 2–epoch combinations. Of these, 66 are rejected as having $i' > 28.0$ mag in the final HUDF stack, leaving only 556 candidates to our magnitude limit. An object where just a single point in the light-curve is deviant would appear as an outlier in 3 out of 6 epoch combinations. This occurs in 25% of the candidates. Another 25% stand out in 2 out of 6 pairs, and 40% stand out in 1 out of 6 pairs (usually indicative of a global rise or decline as a function of time in the light-curve). Of these 556 initial candidates, we find that only 45 out of 4644 show a clear sign of a compact region indicative of a point source, and are devoid of image defects or object splitting issues. These object crowding or splitting issues arise due to the extreme deblending threshold (see § 2.1), which causes unreliable detections and photometric measurements of faint objects in the wings of bright objects. Therefore, we have 45 of 4644 objects that show the signature of AGN variability. The 13 interlopers discussed in § 3 should, in the absence of other information, be evenly distributed amongst the 622 initial candidates, and therefore we expect on the order of 1 out of 45 of our candidates to be a random contaminant. In order to simultaneously show the data from all six possible combinations of the data, the results are plotted in Fig. 3, showing the number of σ by which each object varied in each epoch pair. The colored symbols are for the 45 "best" candidates.

Another 57 objects were found that are "potentially" variable candidates. These are relatively isolated objects with reliable photometry, but show no clear sign of a point source. Since variability has to come from a point-source due to light-travel time considerations, we will ignore these 57 objects for now, but in §6 we will discuss the incompleteness resulting from the finite variability timescales sampled, and from the fraction of non-variable or dust-obscured AGN that have likely been missed al-
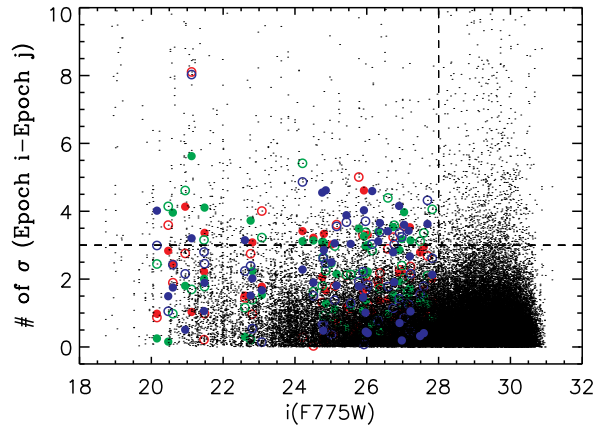

Fig. 3.— Number of σ that each object varies for all six possible combinations of the four different epochs. The "best" sample of 45 variable candidates is shown with colored symbols. These objects were chosen to be $|\Delta mag| \gtrsim 3.0\sigma$ outliers, have $20 \lesssim i' \lesssim 28.0$ mag, and to be unaffected of any local image or weight map structures. Each object appears six times in this plot, and most candidates were seen at $\gtrsim 3\sigma$ in at least two epoch pairs. Significant outliers not plotted in color were almost exclusively due to over-splitting or deblending issues with large objects (mostly occurring in bright, large spiral galaxies), where the enhanced uncertainties in the local (object+sky)-subtraction introduced larger–than–Gaussian flux errors.

together. The four-epoch light-curves for the 45 best candidates are shown in Fig. 4. These light-curve data are tabulated in Table 2, which also specifies the number of epoch pairs where each object was a 3.0σ outlier. For the 45 best candidates, 49% were discovered from a single pair, and 43% in two pairs. Only 5% (2 objects) were found in 3 pairs, which is indicative of a single deviant point in the 4 point light-curve.

In Fig. 5, we show the photometric redshift distribution for all objects with $i' < 28.0$ mag along with that of our best 45 candidates. It is clear that the distribution follows that of the field galaxies, i.e., there is no redshift where faint object variability was most prevalent. To make this more clear, we plot in Fig. 6 the ratio of the $N(z)$ for variables to that of field galaxies, which shows that this fraction is roughly constant at approximately 1% over all redshifts probed in this study. This variability fraction is similar to the 2% found by Sarajedini (2003) in a search for nuclear variability in the HDF–N and the Groth Strip survey.

Interestingly, Straughn et al. (2005) show, in a companion study, that the redshift distribution for "tadpole" galaxies also traces that of the HUDF field galaxies. They argue that these tadpole galaxies are dynamically unrelaxed systems, and therefore in the early stages of merging. The question then arises if tadpole galaxies and objects with point-sources that show signs of AGN variability are drawn from the same population. This question is further addressed in § 6.

5.2. *Other Tests of the Reliability of the Variable Sample*

5.2.1. *Difference Images*

We attempted to also detect the point-source variability using "difference" or "ratio" images between any of the combinations between the two epochs. To do this, we
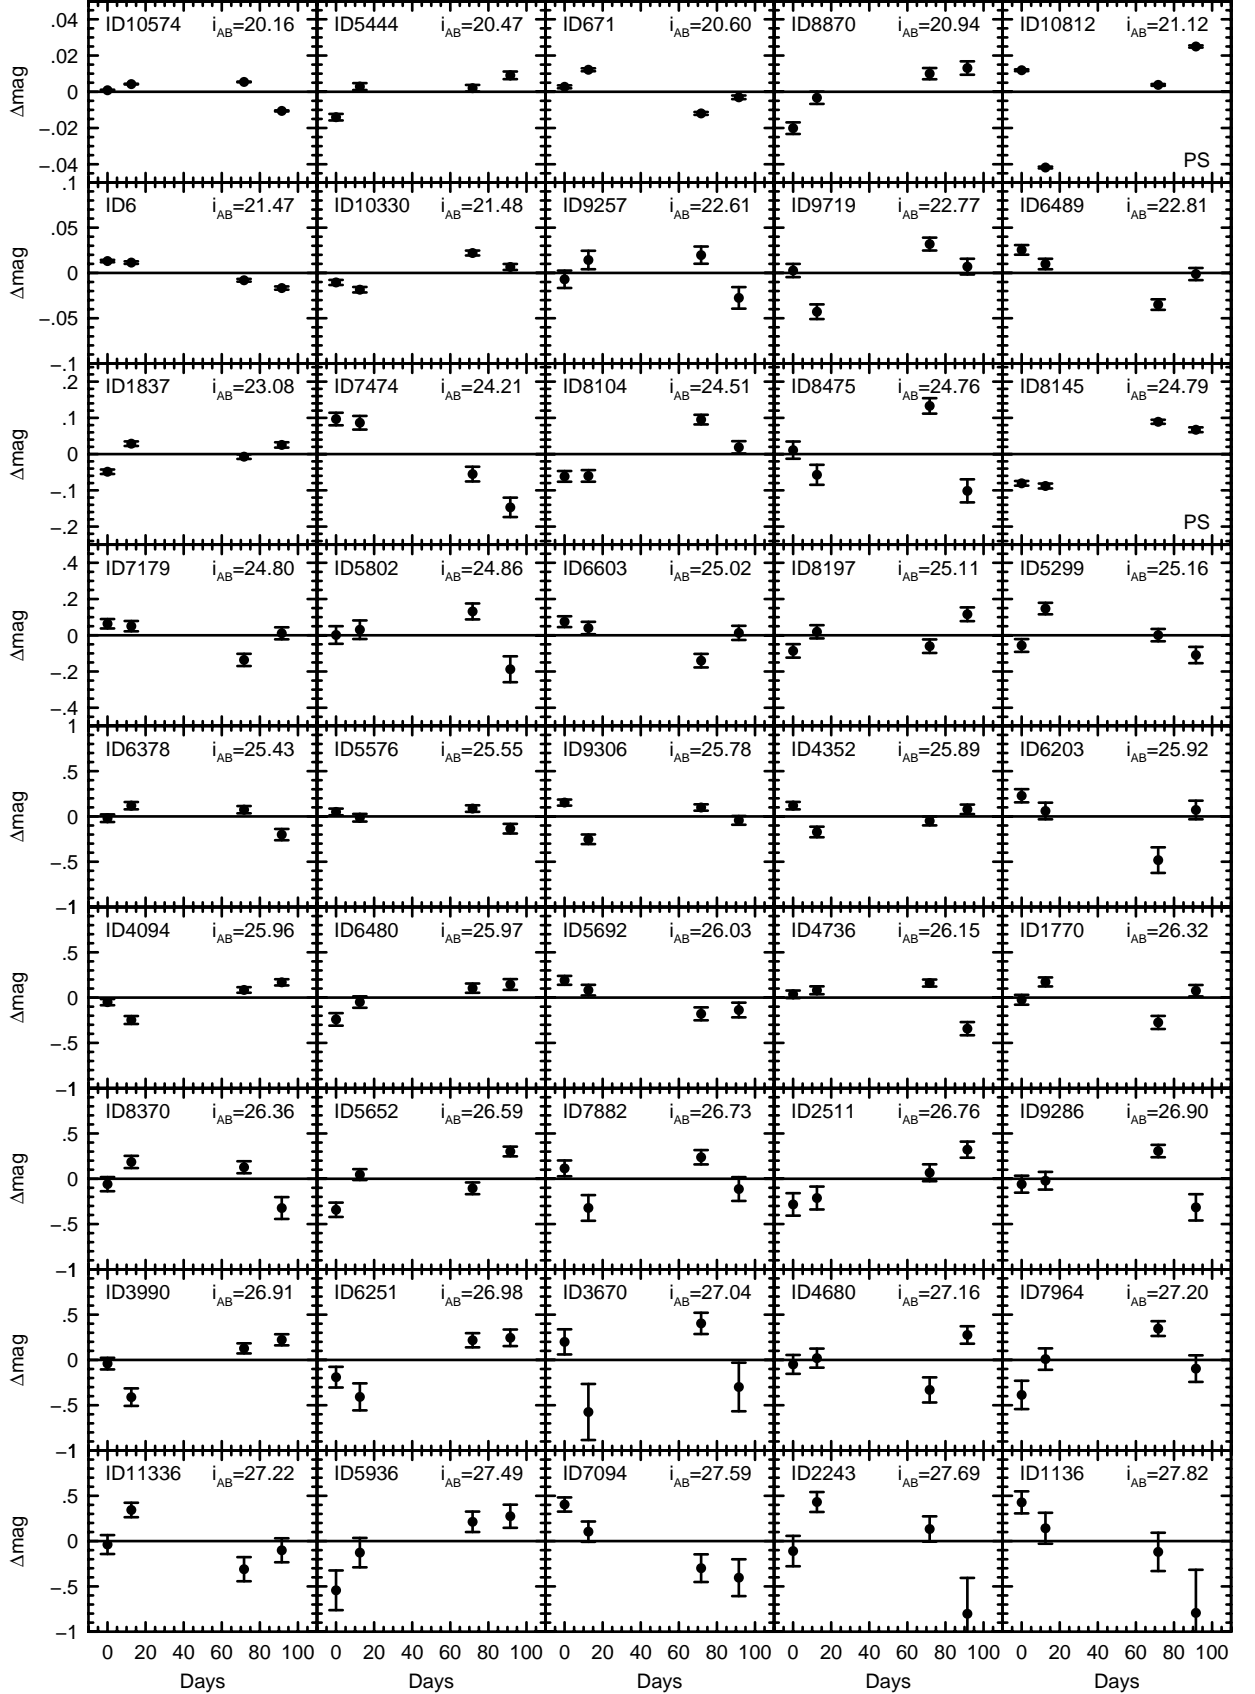
FIG. 4.— Light curves of the 45 best candidates with signs of optical AGN variability. The vertical axes are the change in measured magnitudes plotted in the sense of the average minus individual measurements. The horizontal axis is the number of days since the first measurement. Each panel shows the object ID number in the upper left, and the i'_{AB} magnitude in the upper right. Plots are arranged in order of decreasing flux, and the combined total flux error bars are from *SExtractor*. The two point sources discussed in § 5.2.3 are indicated by "PS."
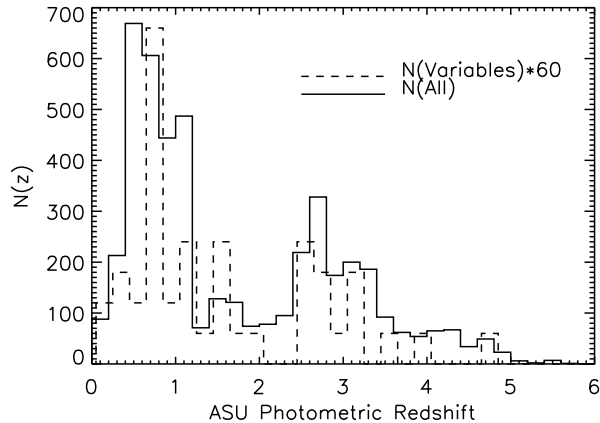
FIG. 5.— Photometric redshift distribution of all HUDF field galaxies (solid line) and for the "best" variable candidates (dashed line) multiplied by 60× for best comparison. Photo-z's computed using *hyperZ* (Bolzonella et al. 2000) and $BVi'z'JH$ HST data for all galaxies with $i' \lesssim 28.0$ mag. The redshift distribution of the variable objects follows that of field galaxies in general.
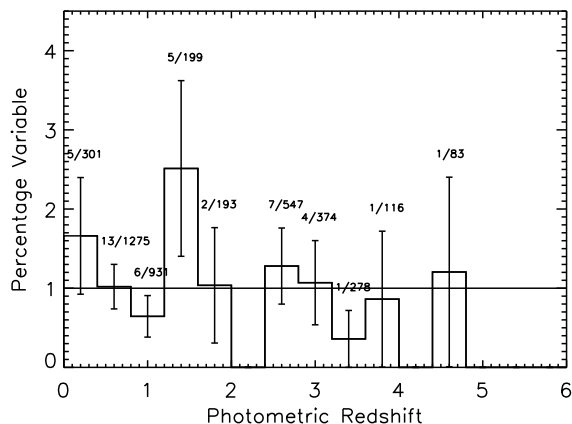

FIG. 6.— Percentage of HUDF objects to $i' \lesssim 28.0$ AB-mag showing variable point sources as a function of photo-z. Within the statistical uncertainties, about 1% of all HUDF galaxies show point source variability over the whole redshift range surveyed ($0 \lesssim z \lesssim 5$).

smoothed the image by a 5×5 and 7×7 box. For each object, a "variance map" is computed over the four epochs as follows. For each pixel, we compute the $n \times n$ box-average for both the image and the weight map. Then we compute the weighted standard deviation over the four epochs for *each* pixel in the box-averaged images. The $n \times n$ box-size was chosen to smooth out most of the PSF breathing and registration issues. Even so, this variance image clearly had peaks in the centers of all bright objects. However, the variable candidates also tended to stand out more than their non-variable neighbors. This method worked best for the brighter objects, clearly verifying 13/18 candidates with $i'_{AB} \lesssim 25$ mag, but only finding ∼50% of the full $i'_{AB} \lesssim 28$ mag sample. While not an optimal method of detection, this ratio-method provided a good way of confirming some of the objects whose total flux changed significantly over the four month period. However, the total flux method remains our primary variable candidate detection method,

since it is most robust against ACS PSF effects.

5.2.2. *Variability Between Two Deeper Epochs*

Given the observing dates listed in Table 1, a natural test is to combine the first and second epochs, and also the third and fourth, giving only two independent HUDF epochs but more widely spaced in time, and with a somewhat higher signal-to-noise ratio for the variable candidate detection. However, we expect fewer variable candidates this way, because there are fewer epochs to compare, and because some of the short-term variability signature is necessarily smoothed over. From this 2–epoch test, 242 candidates are chosen from the error distribution. Next this list is compared to the list of the 45 best candidates from the 4-epoch search, and 12 objects were found in both searches. If the list of the best and possible candidates are combined, there are 102 candidates, and 30 of these are recovered in the two epoch test. This test gives further confidence that at least 30% of our top candidates are truly variable objects, although it does not exclude any of the other ones, since those were found when the shorter time-baselines were included.

A sample of local AGN light-curves can be used in order to assess the expected completeness in temporally condensing our HUDF data from four to two epochs. The best publicly available data is that from the extensive International AGN Watch[1] (hereafter IAGNW; Peterson et al. 2002, and references therein). We used their five IAGNW Seyferts (all with $z < 0.05$), which have B-band light-curves which best match our observed HUDF i'-band data, since the median redshift of our candidates is $z_{med} \simeq 1.5$. Since the rest-frame time sampling depends on the redshifts of the HUDF objects, we re-sampled the IAGNW light-curves for a range of redshifts between $0.5 < z < 5$. The IAGNW data allowed for 30–40 independent 4–epoch light-curves to be created for a given redshift. These light-curves were then scaled to match the observed spread in Δmag as a function of magnitude for our 45 HUDF candidates. A Monte Carlo test was then run on mock catalogs that matched the same magnitude distribution as our 45 HUDF candidates. On average half of the ∼40 light-curves are found when our 4–epoch variability algorithm (using same 3σ selection curves) is applied. The 2–epoch algorithm is then applied to those mock candidates and the number that remain variable candidates turns out to be a decreasing function of assumed redshift. We recover 60–70% for $z \lesssim 1$ and 30–40% for $z \gtrsim 4$. However, this test relies on using 20 templates to simulate 45 objects. A visual inspection of the light-curves in Fig. 4 shows that ideally a larger number of local templates is needed for a fair test. Nonetheless, this limited test that is possible with the available IAGNW data – when taken at face value – may be an indication that as many as one-third of our 45 candidates are potentially false positives, although it is equally likely that the small number of local test templates are not truly representative of the distant HUDF variability sample. Hence, we believe that the reliability of our faint HUDF variability sample is likely larger than 67%, but we will not be able to definitively say for

sure without more local light-curve templates that best match the high-z variable objects.

5.2.3. *Other Wavelengths*

Since the UDF is in the Chandra Deep Field–South (CDFS; Rosati et al. 2002), there exists deep X-ray data in this field, and these data are ideal for detecting AGN. Within the UDF, there are 16 Chandra sources (Koekemoer et al. 2005, in preparation), and we detect four of these as variable in the optical. One of these is a mid-type spiral with $i' = 21.24$ mag, that is a member of a small group of apparently interacting galaxies. Interestingly, this object's total flux varied by less than 1%, but since it was one of the brightest objects in the sample, its variability was clearly detected at $\gtrsim 3.0\sigma$ significance. Two of the others are optical point sources, showing little or no visible host galaxy, with $i' = 21.12$ mag and $i' = 24.79$ mag. They both have measured spectroscopic AGN emission-line redshifts from the GRism ACS Program for Extragalactic Science survey (GRAPES; Pirzkal et al. 2004) at $z \simeq 3$. Our $BVi'z'JH$ photometric redshifts put them at $z = 2.6$ and $z = 2.9$ respectively, which is rather good given that these few brighter AGN are the ones that are most dominated by a non-stellar power-law SED. The fourth candidate is barely detected in any of the ACS bands, and would have been eliminated as a spurious candidate had it not been coincident with the Chandra source, and is clearly a viable object in the VLT K-band image. The detection of 25% of the Chandra sources as optically variable in the HUDF data shows that the method employed here is a reliable way of finding the AGN that are not heavily obscured. Paolillo et al. (2004) find that $> 90\%$ of the studied CDFS X-ray sources with adequate photon statistics show X-ray variability, and a future comparison of the X-ray and optical variability amplitudes could provide insight into the physical mechanism that is responsible for this variability.

5.3. *Possible Sources of Incompleteness in the Variable AGN Sample*

In summary, we found 45 plausibly variable objects with some caveats, and this number may be as high as 102 if the "possible" variable candidates are included. Hence, the true variability fraction on timescales of a few months (rest-frame timescale few weeks to a month) is no more than about 1–2% of all HUDF field galaxies. Except for variable stars such as SNe and novae, this variability is most likely due to an AGN given the timescales and distances involved. Strolger & Riess (2005) found only one moderate redshift SNe in the HUDF, and thus SNe cannot be a significant source of contamination of the present sample.

Two other possible source of incompleteness in the variability study must be addressed first. Non-variable AGN, or AGN that only vary on timescales much longer than 4 months, or optically obscured AGN would not have been detected with the current UV–optical variability method. Sarajedini et al. (2003a) had epochs 5–7 years apart, and found 2% of the objects to be variable. It is thus possible that the currently sampled times-scale shorter than 4 months missed about a factor 2 of all AGN that are only variable on longer time-scales.

A factor of three of the brighter AGN may have been missed, since their UV–optical flux may be obscured by a dust-torus (Barthel 1989). In the AGN unification picture, AGN cones are two-sided and their axes are randomly distributed in the sky, so that an average cone opening angle of ω implies that a fraction 1–$\sin(\omega)$ of all AGN will point in our direction. If $\omega \simeq 45°$ (e.g., Barthel 1989), then every optically detected AGN (QSO) represents 3–4 other bulge-dominated galaxies, whose AGN reflection cone didn't shine in our direction. Hence, their AGN may remain obscured by the dust-torus. Such objects would be visible to Chandra in the X-rays or to Spitzer at mid-IR wavelengths, although the available Chandra and Spitzer data are not nearly deep enough to detect all HUDF objects to $AB = 28$ mag. (Reaching these depths is prohibitive in Chandra and Spitzer integration times, and requires the next generation of X-ray and IR telescopes, such as Generation-X and JWST). At brighter flux limits, Spitzer did indeed recently find a significant fraction of dust-obscured AGN not seen in UV-optical surveys (Urry et al. 2004). In the AGN unification picture, the incompleteness in UV–optically selected samples due to the dust-obscuring torus would be as large as a factor of 3–4 (Treister et al. 2004).

6. DISCUSSION AND CONCLUSIONS

6.1. *Fraction of Variable AGN found*

Interestingly, Straughn et al. (2005) show, in a companion study, that the redshift distribution for "tadpole" galaxies also traces that of the HUDF field galaxies for $0.1 \lesssim z \lesssim 4.5$. They argue that these tadpole galaxies are dynamically unrelaxed systems, and therefore in the early stages of merging. The question then arises if tadpole galaxies and objects with point-sources that show signs of AGN variability are drawn from the same population. At any given redshift, Straughn et al. (2005) find that about 6% of all HUDF galaxies appear as tadpoles, and they conclude that tadpole galaxies are good tracers of the process of galaxy assembly.

Together with the factor of $\gtrsim 2$ incompleteness in the HUDF variability sample due to the limited time-baseline sampled thus far, the actual fraction of weak AGN present in these dynamically young galaxies may be a factor of $\gtrsim 6$–$8\times$ larger than the 1% variable AGN fraction found through variability in the HUDF. Hence, perhaps as many as $\gtrsim 6$–8% of all field galaxies may host weak AGN, only $\sim 1\%$ of which we found here, and another $\gtrsim 1\%$ could have been found if longer time-baseline had been available. The other factor of 3–4$\times$ of AGN are likely missing because they are optically obscured, requiring the next generation of X-ray and IR telescopes.

6.2. *AGN feeding timescales compared to galaxy merger timescales*

We now consider if the current variability dataset can constrain the merging rate of SMBHs, assuming that SMBHs formed during hierarchical mergers of galaxies containing less massive SMBHs, and if signatures of galaxy mergers could be related to variable AGNs. Put another way, can signatures of galaxy merging be related to our variable AGN candidates?

A closer inspection of our data reveals that only 1 or perhaps 2 of the variable candidates resemble

the tadpole galaxies of Straughn et al. (2005). Recent state-of-the-art hierarchical models have suggested that the AGN-phase *only* occurs in the later stages of a galaxy merger, and well *after* it appears in the tadpole phase (Springel, Di Matteo, & Hernquist 2005; di Matteo, Springel, & Hernquist 2005). These models also predict that the AGN will likely only be visible *well after* ($\gtrsim 1$ Gyr) the merger induced star-formation has died down, implying that the fraction of dynamically young tadpoles that are expected to already show active AGN properties is relatively small. The small overlap between the two populations that we observe does thus not exclude the possibility that faint object variability is tracing the growth of SMBHs. However, this SMBH growth can only be constrained indirectly from the variability data, and a more detailed discussion of the connection between mergers and AGN activity is given in Straughn et al. (2005) and references therein.

Given the importance of understanding the growth and origin of SMBHs, several lessons can be learned from this work in order to better design future studies of this type. The time-spacing of the HUDF observations –although as good as possible given scheduling constraints–was not ideal for this type of study. It is critical to re-visit the HUDF with HST, with the observations optimized for a faint-object variability study, including covering time-scales of a few years. It is also essential to plan deeper surveys at longer wavelengths with the JWST. The JWST photometric *and* PSF stability are crucial in this regard, as many of our HUDF objects show significant variability of less than a few percent in flux. Also, a limiting factor in our results is the breathing of the PSF, along with image noise due to correcting the geometric distortion of ACS and slight variations of the PSF across the field. Therefore, JWST must have design specifications that are capable of meeting these requirements in order to permit faint object variability studies to be done.

Facility: HST(ACS).

REFERENCES

Barthel, P. D. 1994, ApJ, 336, 606
Beckwith, S. et al. 2005, AJ, in prep.
Bertin, E. & Arnouts 1996, A&AS, 117,393
Bolzonella, M., Miralles, J.-M., & Pelló, R. 2000, A&A, 363,476
Bouwens, R. J, et al. 2004, ApJ, 616, 79
Bruzual, G. & Charlot, S. 1993 ApJ, 405, 538
Coleman, G. D., Wu, C.-C., & Weedman, D. W. 1980, ApJS, 43, 393
di Matteo, T., Springel, V., & Hernquist, L. 2005, Nature, 433, 604
Ferrarese, L. & Ford, H.C. 2004, Space Sci. Rev., 116, 523
Heymans, C., et al. 2005, MNRAS, 361, 160
Koekemoer, A. M., Fruchter,A. S., Hook, R., & Hack, W. 2002, in Proceedings of the 2002 HST Calibration Workshop, ed. S. Arribas, A. Koekemoer, and B. Whitmore (Baltimore: Space Telescope Science Institute), 337
Kogut, A., et al. 2003, ApJS, 148, 161
Kormendy, J. & Richstone, D. 1995, ARA&A, 33, 581
Kormendy, J. & Gebhardt, K. 2001, in AIP Conf. Proc. 586: 20th Texas Symposium on Relativistic Astrophysics, 363
Madau, P., & Rees, M. 2001, ApJ, 551, L27

Magorrian, J., Tremaine, S., & Richstone, D. 1998, AJ,115, 2285
Paltani, S., & Courvoisier, T. J.-L. 1994, A&A, 291, 74
Paolillo, M., Schreier, E. J., Giacconi, R., Koekemoer, A. M., & Grogin, N. A. 2004, ApJ, 611, 93
Pirzkal, N., et al. 2004, ApJS, 154, 501
Peterson, B. M., et al. 2002, ApJ, 581, 197
Rosati, P., et al. 2002, ApJ, 566, 667
Sarajedini, V. L., Gilliland, R. L. & Kasm, C. 2003, ApJ, 599, 173
Sarajedini, V. L., 2003, MmSAI, 74, 957
Silk, J., & Rees, M. 1998, A&A, 331, L1
Springel,V., Di Matteo, T., & Hernquist, L. 2005, ApJ, 620, 79
Straughn, A. N., Cohen, S. H., Ryan, R. E., Hathi, N. P., & Windhorst, R. A. 2005, this volume
Strolger, L.-G., & Riess, A. G. 2005, preprint (astro-ph/0503093)
Thompson, R. I., et al. 2005, AJ, 130, 1
Treister, E., et al. 2004, ApJ, 616, 123
Urry, C. M., et al. 2004, BAAS, 205.6209
Yan, H., & Windhorst, R. A. 2004, ApJ, 612, 93

TABLE 2
Best 45 Variability Candidates

ID	RA (J2000)	Dec (J2000)	i_{AB}[a]	$\langle i_{AB}\rangle$[b]	$\Delta mag1$[c]	$\Delta mag2$[c]	$\Delta mag3$[c]	$\Delta mag4$[c]	#[d]	Notes[e]
	(h m s)	(° ′ ″)	(mag)	(mag)	(mmag)	(mmag)	(mmag)	(mmag)		
10574	03 32 37.92	-27 46 09.1	20.16	20.16	1± 0	4± 0	5± 0	-11± 0	1	⋯
5444	03 32 44.98	-27 47 36.9	20.47	20.49	-14± 2	3± 2	2± 2	9± 2	2	⋯
671	03 32 41.09	-27 48 53.0	20.60	20.61	3± 1	12± 1	-12± 1	-3± 1	1	Two
8870	03 32 36.67	-27 46 31.1	20.94	20.97	-20± 3	-3± 3	10± 3	13± 4	2	Two, CXO
10812	03 32 39.09	-27 46 01.8	21.12	21.12	12± 1	-42± 1	4± 1	25± 1	4	Two, CXO, PS
6	03 32 39.54	-27 49 28.4	21.47	21.48	13± 1	11± 1	-8± 2	-17± 2	1	Two
10330	03 32 37.19	-27 46 08.1	21.48	21.51	-11± 3	-19± 3	22± 3	7± 3	2	Two
9257	03 32 44.28	-27 46 42.3	22.61	22.63	-7± 10	14± 10	20± 10	-28± 12	1	⋯
9719	03 32 36.43	-27 46 32.6	22.77	22.82	3± 7	-43± 8	32± 7	7± 9	1	⋯
6489	03 32 44.78	-27 47 24.8	22.81	22.81	25± 5	10± 6	-35± 6	-1± 7	1	⋯
1837	03 32 33.12	-27 48 29.6	23.08	23.12	-49± 6	29± 6	-7± 6	25± 7	2	⋯
7474	03 32 31.51	-27 47 12.3	24.21	24.37	97± 17	87± 19	-55± 20	-147± 27	4	Two
8104	03 32 42.86	-27 47 02.7	24.51	24.53	-62± 15	-60± 16	95± 13	19± 17	2	CXO?
8475	03 32 38.99	-27 46 56.7	24.76	24.95	11± 24	-57± 28	133± 21	-101± 32	2	⋯
8145	03 32 42.83	-27 47 02.5	24.79	24.78	-81± 6	-88± 7	89± 6	67± 7	2	Two, CXO?, PS
7179	03 32 30.17	-27 47 16.9	24.80	24.87	64± 26	50± 29	-136± 34	11± 33	2	⋯
5802	03 32 37.49	-27 47 31.4	24.86	25.46	2± 49	31± 51	132± 44	-188± 72	1	⋯
6603	03 32 48.21	-27 47 24.1	25.02	25.20	75± 30	40± 34	-140± 37	13± 39	1	⋯
8197	03 32 31.83	-27 47 02.9	25.11	25.17	-86± 37	19± 37	-60± 37	116± 38	1	⋯
5299	03 32 36.05	-27 47 37.8	25.16	25.18	-56± 35	147± 32	1± 34	-109± 45	2	⋯
6378	03 32 36.20	-27 47 26.2	25.43	25.47	-20± 42	119± 40	74± 40	-200± 62	2	⋯
5576	03 32 42.23	-27 47 33.4	25.55	25.56	50± 37	-14± 42	87± 36	-135± 53	1	⋯
9306	03 32 41.09	-27 46 42.4	25.78	25.77	153± 33	-252± 53	99± 35	-43± 49	2	⋯
4352	03 32 33.93	-27 47 49.9	25.89	25.94	119± 41	-172± 59	-50± 49	78± 52	1	⋯
6203	03 32 37.00	-27 47 26.3	25.92	26.59	228± 73	60± 92	-482±141	71±103	3	⋯
4094	03 32 33.11	-27 47 52.1	25.96	25.95	-50± 34	-247± 44	84± 31	169± 34	1	Two
6480	03 32 41.40	-27 47 23.9	25.97	26.08	-241± 69	-49± 63	104± 51	145± 59	1	⋯
5692	03 32 44.97	-27 47 33.6	26.03	26.08	191± 49	82± 59	-179± 71	-137± 81	1	Two
4736	03 32 43.25	-27 47 44.0	26.15	26.27	36± 42	81± 44	161± 39	-343± 73	2	⋯
1770	03 32 35.75	-27 48 31.5	26.32	26.30	-24± 54	173± 50	-274± 72	76± 63	2	⋯
8370	03 32 41.97	-27 46 58.1	26.36	26.48	-59± 78	186± 67	127± 66	-323±121	2	⋯
5652	03 32 42.05	-27 47 33.1	26.59	26.61	-342± 79	46± 60	-105± 65	301± 54	2	⋯
7882	03 32 45.76	-27 47 05.7	26.73	26.67	115± 87	-322±142	238± 79	-113±131	2	⋯
2511	03 32 44.36	-27 48 16.2	26.76	26.90	-283±124	-213±126	66± 93	321± 89	2	Two
9286	03 32 45.23	-27 46 42.5	26.90	26.95	-60± 94	-22± 98	306± 68	-316±145	1	⋯
3990	03 32 33.14	-27 47 52.7	26.91	26.89	-41± 64	-411± 97	127± 56	222± 61	1	⋯
6251	03 32 41.05	-27 47 25.7	26.98	27.03	-190±114	-408±149	217± 78	244± 91	2	Two
3670	03 32 35.94	-27 47 58.0	27.04	27.52	199±139	-575±309	403±117	-299±268	3	⋯
4680	03 32 38.62	-27 47 44.7	27.16	27.24	-49±104	20±105	-331±139	275± 96	1	⋯
7964	03 32 44.27	-27 47 03.6	27.20	27.25	-386±156	9±118	346± 82	-96±147	1	⋯
11336	03 32 39.68	-27 45 46.0	27.22	27.25	-38±104	344± 80	-310±133	-101±132	1	⋯
5936	03 32 45.20	-27 47 29.1	27.49	27.61	-543±219	-127±162	213±113	274±128	1	⋯
7094	03 32 44.26	-27 47 14.9	27.59	27.58	404± 78	104±112	-299±153	-404±203	1	Two
2243	03 32 35.93	-27 48 20.5	27.69	27.69	-110±168	432±110	133±140	-802±396	2	⋯
1136	03 32 37.58	-27 48 46.3	27.82	27.90	428±122	142±171	-119±211	-793±476	1	⋯

[a]Measured magnitude from total UDF stack.
[b]Magnitude of average flux over each of the four epochs.
[c]Difference between average magnitude and magnitude from Epoch N in milli-mags.
[d]Number of epoch pairs where variability detected (max is 6).
[e]Two–also detected in two epoch test; CXO–located at position of CXO source; PS–point source (see § 5.2.3).